

**Errata for the
PCI Express® Base Specification
Revision 3.1,
Single Root I/O Virtualization and
Sharing Revision 1.1,
Address Translation and Sharing
Revision 1.1, and M.2 Specification
Revision 1.0**

December 13, 2017



**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

REVISION	REVISION HISTORY	DATE
1.0	First Release. Includes B208, A210, B211, B214 to B216, A217, B218 to B220, B224, B228 to B237, B242, B244, B246, B248 to B250, B252 to B264	2015-09-18
2	Final Release against Base Revision 3.1a. Subsequent Errata will be against Base Revision 4.0	2017-12-13

PCI-SIG® disclaims all warranties and liability for the use of this document and the information contained herein and assumes no responsibility for any errors that may appear in this document, nor does PCI-SIG make a commitment to update the information contained herein.

Contact the PCI-SIG office to obtain the latest revision of this specification.

Questions regarding the PCI Express Base Specification or membership in PCI-SIG may be forwarded to:

Membership Services

www.pcisig.com

E-mail: administration@pcisig.com

Phone: 503-619-0569

Fax: 503-644-6708

Technical Support

techsupp@pcisig.com

DISCLAIMER

This PCI Express Base Specification Errata document is provided “as is” with no warranties whatsoever, including any warranty of merchantability, noninfringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. PCI-SIG disclaims all liability for infringement of proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

PCI, PCI Express, PCIe, and PCI-SIG are trademarks or registered trademarks of PCI-SIG.

All other product names are trademarks, registered trademarks, or servicemarks of their respective owners.

Copyright © 2011-2017 PCI-SIG

Contents

B1	Separate RefClk with Independent SSC (SRIS) Equation 4.3.9	5
B208	Spread Spectrum Modulation Slew Rate	5
A210	M-PCIe Configuration.Update	8
B211	Root Complex Extended Capability Headers	9
B214	Inference of Electrical Idle	10
B215	ACS Source Validation and Bus Number 0	10
B216	Root Complex Integrated Endpoint Terminology	11
A217	L1 PM Substates: L1.2 Entry Conditions.....	11
B218	Root Complex and Routing Element Terms	12
B244	Device Terminology	12
B219	L1 PM Substates: Electrical Idle Terminology	12
B220	LTSSM Disabled to Detect Transition	13
B224	Malformed TLPs and Disabled VCs	13
B228	Link Speed Management	14
B229	Corrupt Received DLLP	15
B230	Loopback	15
B231	Flow Control Information Tracked by Receiver	16
B233	M.2 Specification: Reference to Link Capabilities register	16
B234	WAKE# and OBFF	17
B235	Readiness Time Reporting example times.....	17
B236	SR-IOV: Changing NumVFs or ARI Capable Hierarchy	18
B237	SR-IOV: Bus Numbering inside a Root Complex	18
B242	ATS: Terminology: PRG Response PASID Required	18
B246	Device Serial Number	19
B248	ATS: Page Request Interface and Dirty Bits	19
B249	Access Control Services for SR-IOV	21
B250	Enhanced Allocation: Address Maps.....	23
B252	Latency Tolerance Reporting (LTR)	23
B253	LCRC and Sequence Number Rules (TLP Transmitter)	24
B254	PCI Express Capability Structure	24
B255	Enhanced Allocation: Reserved Encodings	25
B256	Enhanced Allocation: Virtual Function BEI Ranges.....	25
B257	Lane Equalization Control Register Nomenclature.....	25
B258	Timing Recommendations for Latency Tolerance Reporting	27
B259	Completion Timeout Mechanism.....	29
B260	Slot Power Limit	29
B263	AtomicOp Completion Rules	29
B264	PASID Completions	30
B265	SR-IOV: ARI Capable Hierarchy	31

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

C266	LNR Control Register Offset	31
C267	Resizable BAR Control Register Name Usage	31
C268	Resizable BAR Capabilities Register Name Usage	31
B269	DC Balance Tracking	32
A270	ECR and ECN Process	32
A271	Expanded Resizable BAR: BAR Size field	32
B274	Advisory Non-Fatal.....	33
C275	MFVC term usage	33
B276	DRS and FRS Message Subtype.....	34
B277	ATS: Translation Completion	35
B278	ATS: Translation Completion Length	35
B279	ATS: Completions with Multiple Translations	36
B280	ATS: Translation Completion Lower Address Field.....	36
B281	Completer breaking up Responses	36
B282	Reserved bits in SKP OS Generation Vector	38
B283	Enhanced Allocation ECN: Enhanced Allocation Capability.....	39
B284	VF Resizable Bar: VF BAR Size	42
B285	PASID TLP Prefix.....	43
B286	Multi-Function Arbitration Model Example.....	45
B288	Flow Control for Nullified TLPs.....	45
B289	Clarification of HwInit	46
B292	Capturing Device Number.....	46
B293	DPC Header Log.....	47
B294	Completion Handling Rules for ID Based Ordering	47
B295	Downstream Port Containment Clarifications.....	47
B297	ATS: Clarifications.....	49
B299	PCI Local Bus and SR-IOV: Vendor, Device, and Revision IDs.....	50
B300	PME Generation.....	52
B301	PCI Express Capability Structure	54
B303	Hierarchy ID ECN: Hierarchy ID Message	55

Note on Errata Titles and Errata Numbering

Red text in the errata title indicates the affected document(s). Absence of red text means PCI Express Base Revision 3.1a.

Errata numbering is arbitrary and reflects numbers used during the development process.

B1 Separate RefClk with Independent SSC (SRIS) Equation 4.3.9

Note: *This erratum was incorporated into PCI Express Base 3.1a.
This is actually the only difference between PCI Express Base 3.1 and 3.1a.*

In Section 4.3.8.5, correct Equation 4.3.9 as follows:

$$H(s) = \frac{s^2}{s^2 + sA + B} \times \frac{s^2 + 2\zeta_2\omega_0s + \omega_0^2}{s^2 + 2\zeta_1\omega_0s + \omega_0^2}$$

where:

$$\begin{aligned}\zeta_1 &= \frac{1}{\sqrt{2}} \\ \zeta_2 &= 1 \\ \omega_0 &= 10^7 \times 2\pi \\ A &= 10^7 \times 2\pi \\ B &= 2.2 \times 10^2 \times (2\pi)^2\end{aligned}$$

B208 Spread Spectrum Modulation Slew Rate

In Section 4.3.7.3.3, page 427, make the following changes:

Table 4-32: Refclk Parameters for Common Refclk Rx Architecture at 5.0 GT/s

Symbol	Description	Limits		Units	Note
		Min	Max		
T _{REFCLK-HF-RMS}	> 1.5 MHz to Nyquist RMS jitter after applying Equation 4.3.3		3.1	ps RMS	1
T _{REFCLK-SSC-RES}	SSC residual		75	ps	1
T _{REFCLK-LF-RMS}	10 kHz - 1.5 MHz RMS jitter		3.0	ps RMS	2
T _{SSC-FREQ-DEVIATION}	SSC deviation		+0.0/-0.5	%	
T _{SSC-MAX-PERIOD-SLEW}	Maximum <u>SSC df/dt rate of change of the clock frequency</u>		0.75 <u>1250</u>	ps/UI <u>ppm/μs</u>	3

Notes:

1. T_{REFCLK-HF-RMS} is measured at the far end of the test circuit illustrated in Figure 4-89 after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
2. T_{REFCLK-SSC-RES} and T_{REFCLK-LF-RMS} are measured after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

3. Defined for a worst case SSC modulation profile such as Lexmark.

In Section 4.3.7.3.4, page 428, line 12, make the following changes:

In the data clocked Rx architecture, the amount of Refclk jitter propagated depends on the maximum PLL bandwidth (16 MHz, 3 dB of peaking). It is also the case that the Rx CDR must track the entirety of SSC, (20 ns) and the CDR must be capable of tracking SSC at a maximum slew rate corresponding to the largest ~~df/dt~~ clock frequency change rate SSC modulation profiles.

In Section 4.3.7.3.5, page 429, make the following changes:

Table 4-34: Refclk Parameters for Data Clocked Rx Architecture

Symbol	Description	Limits		Units	Note
		Min	Max		
T _{REFCLK-HF-RMS}	≥ 1.5 MHz to Nyquist RMS jitter after applying Equation 4.3.5		4.0	ps RMS	1
T _{REFCLK-SSC-RES}	Full SSC modulation corresponding to +0 – 0.5%		20	ns	1
T _{REFCLK-LF-RMS}	10 kHz - 1.5 MHz RMS jitter		7.5	ps RMS	2
T _{SSC-FREQ-DEVIATION}	SSC deviation		+0.0/-0.5	%	
T _{SSC-MAX-PERIOD-SLEW}	Max SSC df/dt <u>Maximum rate of change of the clock frequency</u>		0.75 <u>1250</u>	ps/UI <u>ppm/μs</u>	3

Notes:

1. T_{REFCLK-HF-RMS} is measured at the far end of the test circuit illustrated in Figure 4-92 after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
2. T_{REFCLK-SSC-RES} and T_{REFCLK-LF-RMS} are measured after the filter function defined in Table 4-29 for Common Refclk Rx for >1.5 MHz jitter components has been applied.
3. Defined for a worst case SSC modulation profile such as Lexmark.

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

In Section 4.3.7.5, page 433, add note reference for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-35: Parameters for Separate Refclk With Independent SSC Architecture (SRIS) at 5.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F_{REFCLK}	Refclk frequency	99.97	100.03	MHz
$T_{REFCLK-RMS-SRIS}$	RMS Refclk jitter for Separate Refclk Independent SSC architecture ¹		2.0	ps RMS
F_{SSC}	SSC frequency range	30	33	kHz
$T_{SSC-FREQ-DEVIATION}$	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying PLL transfer function (16 MHz, 2 dB) and CDR jitter transfer function defined in Figure 4-95.
2. The maximum rate of change of the clock frequency is 1250 ppm/ μ s.

In Section 4.3.8.2, page 436, line 4, make the following changes:

Notes:

1. Before application of SSC.
2. ~~It is sufficient to define SSC deviation only without specifying anything about the shape of the modulation envelope. Envelopes with very large df/dt will fail the $T_{REFCLK-RMS-CC}$ parameter.~~ The maximum rate of change of the clock frequency is 1250 ppm/ μ s.
3. The Refclk jitter is measured after applying the jitter filtering function defined in Figure 4-96.

In Section 4.3.8.3, page 438, add note for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-38: Parameters for Data Clocked Rx Architecture at 8.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F_{REFCLK}	Refclk frequency	99.97	100.03	MHz
$T_{REFCLK-RMS-SRIS}$	RMS Refclk jitter for data clocked Rx architecture ¹		1.0	ps RMS
F_{SSC}	SSC frequency range	30	33	kHz
$T_{SSC-FREQ-DEVIATION}$	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying jitter filtering function defined in Figure 4-97.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

2. [The maximum rate of change of the clock frequency is 1250 ppm/μs.](#)

In Section 4.3.8.5, page 440, add note reference for $T_{SSC-FREQ-DEVIATION}$ as follows:

Table 4-39: Parameters for Separate Refclk With Independent SSC (SRIS) Architecture at 8.0 GT/s

Symbol	Description	Limits		Units
		Min	Max	
F_{REFCLK}	Refclk frequency	99.97	100.03	MHz
$T_{REFCLK-RMS-SRIS}$	RMS Refclk jitter for Separate Refclk Independent SSC architecture ¹		0.5	ps RMS
F_{SSC}	SSC frequency range	30	33	kHz
$T_{SSC-FREQ-DEVIATION}$	SSC deviation ²		+0.0/-0.5	%

Notes:

1. The Refclk jitter is measured after applying PLL transfer function (4 MHz, 2 dB) and CDR jitter transfer function defined in Figure 4-98.
2. The maximum rate of change of the clock frequency is 1250 ppm/μs.
3. The maximum rate of change of the clock phase is 3 ns/μs.

A210 M-PCIe Configuration.Update

In Section 8.4.9.3.3, page 963, line 15, make the following changes:

8.4.9.3.3. Configuration.Update

- ☐ On entry, all the TX-LANES and RX-LANES are in HIBERN8 state.
- ☐ The recovery_to_configuration variable must be set to 0b.
- ☐ The Configured TX-LANES must exit to STALL after residing in this state for THIBERN8.
- ☐ The next state is Configuration.Confirm if all Configured TX-LANES have resided in STALL for at least TACTIVATE, where TACTIVATE is measured after detecting DIF-N on all the RX-LANES that are part of the new LINK width.
- ☐ [Otherwise, the next state is Configuration.ExitToDetect after a 2 ms timeout.](#)

In Section 8.4.9.3.7, page 965, line 15, make the following changes:

8.4.9.3.7. Configuration.ExitToDetect

- ☐ The following steps must be followed in sequence:

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

- ~~All the~~ Configured TX-LANes that are in an HS-BURST must send at least one EIOS, terminate the HS-BURST and enter STALL.
- All Configured LANes initiate the process to enter HIBERN8 as described in Section 8.4.10.
- Local Reset must be asserted to all M-PHY MODULES associated with the LINK.
- The next state is Detect

B211 Root Complex Extended Capability Headers

In Section 7.13.1, Table 7-56, page 789, line 10, make the following changes:

Table 7-56: Root Complex Link Declaration Extended Capability Header

Bit Location	Register Description	Attributes
15:0	PCI Express Extended Capability ID – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability. <u>The</u> Extended Capability ID for the <u>Root Complex</u> Link Declaration <u>Extended</u> Capability is 0005h.	RO

In Section 7.13.2, page 790, line 1, make the following changes:

The Element Self Description register provides information about the Root Complex element containing the Root Complex Link Declaration Capability.

In Section 7.14.1, Table 7-60, page 796, line 6, make the following changes:

Table 7-60: Root Complex Internal Link Control Extended Capability Header

Bit Location	Register Description	Attributes
15:0	PCI Express Extended Capability ID – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability. <u>The</u> Extended Capability ID for the <u>Root Complex Internal Link Control Extended</u> Link Declaration Capability is 0006h.	RO

B214 Inference of Electrical Idle

In Section 4.2.4.3, Page 257, line 4 make the following changes:

4.2.4.3. Inferring Electrical Idle

...

IMPLEMENTATION NOTE

Inference of Electrical Idle

In the L0 state, ~~some number of one or more~~ Flow Control Update DLLPs are expected to be received in a 128 μ s window. Also in L0, ~~a one or more~~ SKP Ordered Sets ~~is~~ are expected to be received in a 128 μ s window. ~~on average every 1538 Symbols in 2.5 GT/s and 5.0 GT/s data rates and every 375 Blocks in 8.0 GT/s data rate.~~ As a simplification, it is permitted to use either one (or both) of these indicators to infer Electrical Idle. Hence, the absence of ~~either~~ a Flow Control Update DLLP and/or ~~alternatively~~ a SKP Ordered Set in any 128 μ s window can be inferred as Electrical Idle.

B215 ACS Source Validation and Bus Number 0

In Section 6.12.1.1, page 588, line 12, make the following changes:

6.12.1.1. ACS Downstream Ports

...

- ☐ ACS Source Validation: must be implemented.

When enabled, the Downstream Port tests the Bus Number from the Requester ID of each Upstream Request received by the Port to determine if it is within the Bus Number “aperture” of the Port – the inclusive range specified by the Secondary Bus Number register and the Subordinate Bus Number register.

If the Bus Number from the Requester ID of the Request is not within this aperture, this is a reported error (ACS Violation) associated with the Receiving Port (see Section 6.12.4.)

Completions are never affected by ACS Source Validation.

IMPLEMENTATION NOTE:

Upstream Messages and ACS Source Validation

Functions are permitted to transmit Upstream Messages before they have been assigned a Bus Number. Such messages will have a Requester ID with a Bus Number of 00h. If the Downstream Port has ACS Source Validation enabled, these Messages (see Section 2.2.8.6 and Section 6.23.1) will likely be detected as an ACS Violation error.

- ☐ ACS Translation Blocking: must be implemented.

...

In Section 6.23.1, page 651, line 20, make the following changes:

6.23.1. Device Readiness Status (DRS)

...

Additional requirements for Root Ports and Switch Downstream Ports include:

- ☐ Implementation of the DRS Message Received bit, which indicates receipt of a DRS Message

IMPLEMENTATION NOTE:

DRS Messages and ACS Source Validation

Functions are permitted to transmit DRS Messages before they have been assigned a Bus Number. Such messages will have a Requester ID with a Bus Number of 00h. If the Downstream Port has ACS Source Validation enabled, these Messages (see Section 6.12.1.1) will likely be detected as an ACS Violation error.

B216 Root Complex Integrated Endpoint Terminology

*Replace all 13 occurrences of '**RCIE**' with 'RCiEP'.*

*Replace all 3 occurrences of '**RCIEs**' with 'RCiEPs'.*

In the Terms in Acronyms Section, page 42, line 15, add the following definition:

RCiEP Root Complex Integrated End Point.

A217 L1 PM Substates: L1.2 Entry Conditions

In Section 5.5.1, page 488, line 15, make the following changes:

5.5.1. Entry conditions for L1 PM Substates and L1.0 Requirements

...

- ☐ When in ASPM L1.0 and the ASPM L1.2 Enable bit is Set, the L1.2 substate must be entered when CLKREQ# is de-asserted and all of the following conditions are true:
 - ~~†~~The reported snooped LTR value last sent or received by this Port is greater than or equal to the value set by the LTR_L1.2_THRESHOLD Value and Scale fields, or there is no snoop service latency requirement.
 - ~~†~~The reported non-snooped LTR value last sent or received by this Port is greater than or equal to the value set by the LTR_L1.2_THRESHOLD Value and Scale fields, or there is no non-snoop service latency requirement.

B218 Root Complex and Routing Element Terms

In the Terms and Acronyms Section, page 42, line 36, make the following changes:

Terms and Acronyms

...

Root Complex, RC	A defined System Element that includes at least one Host Bridge, Root Port, or Root Complex Integrated Endpoint . zero or more Host Bridges, zero or more Root Complex Integrated Endpoints, zero or more Root Complex Event Collectors, and one or more Root Ports.
Root Complex Component	A logical aggregation of Root Ports, Root Complex Register Blocks, Root Complex Integrated Endpoints, and Root Complex Event Collectors.
Root Port	A PCI Express Port on a Root Complex that maps a portion of a Hierarchy through an associated virtual PCI-PCI Bridge.
RP PIO	Root Port Programmed I/O. See Section 6.2.10.3.
Routing Element	A term referring to a Root Complex, Switch, or Bridge in regard to its ability to route, multicast, or block TLPs.
RRAP	Remote Register Access Protocol (RRAP) is ...

B244 Device Terminology

In the Terms and Acronyms Section, page 38, line 15, make the following changes:

Device	A collection of one or more Functions within a single Hierarchy identified by common Bus Number and Device Number. An SR-IOV Device may have additional Functions accessed via additional Bus Numbers and/or Device Numbers configured through one or more SR-IOV Capability structures.
--------	--

B219 L1 PM Substates: Electrical Idle Terminology

In Section 5.5, page 485, line 1, make the following changes:

5.5. L1 PM Substates

...

☐ L1.1 substate

...

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

- The Upstream and Downstream Ports are not required to be enabled to detect Electrical Idle [exit](#).

...

❑ L1.2 substate

...

- The Upstream and Downstream Ports are not required to be enabled to detect ~~electrical idle~~ ~~(EI)~~ [Electrical Idle exit](#).

B220 LTSSM Disabled to Detect Transition

In Section 4.2.6.9, page 326, line 17, make the following changes:

4.2.6.9. Disabled

...

- ❑ If an EIOS was transmitted (one if the current Link speed is 2.5 GT/s or 8.0 GT/s and two consecutive ones if the current Link speed is 5.0 GT/s) and an EIOS was received on any Lane (even while transmitting TS1 with the Disable Link bit asserted), then:
 - LinkUp = 0b (False)
 - At this point, the Lanes are considered Disabled.
 - For Upstream Ports: All Receivers must meet the Z_{RX-DC} specification for 2.5 GT/s within 1 ms (see Table 4-24).
 - For Upstream Ports: The next state is Detect when Electrical Idle ~~E~~[exit](#) is detected [on at least one Lane](#) ~~at the Receiver~~.

B224 Malformed TLPs and Disabled VCs

In Section 2.3, page 123, line 22, make the following changes:

2.3. Handling of Received TLPs

...

- Received Malformed TLPs that are ambiguous with respect to which buffer to release or are mapped to an uninitialized [or disabled](#) Virtual Channel must be discarded without updating Receiver Flow Control information.

In Section 5.3.2.1, page 454, line 10, make the following changes:

5.3.2.1. Entry into the L1 State

...

3. The Downstream component must then wait until it accumulates at least the minimum number of credits required to send the largest possible packet for any FC type [for all enabled VCs](#) (if it does not already have such credits). All Transaction Layer TLP scheduling is then suspended.

In Section 5.4.1.2.1, page 472, line 13, make the following changes:

5.4.1.2.1. Entry into the L1 State

...

- ❑ The Downstream component must not initiate ASPM L1 entry until it accumulates at least the minimum number of credits required to send the largest possible packet for any FC type [for all enabled VCs](#).

B228 Link Speed Management

In Section 6.11, page 585, line 7, make the following changes:

6.11. Link Speed Management

...

The Target Link Speed field in the Link Control 2 register [in the Downstream Port](#) sets the upper bound for the Link speed. Except as described below, the Upstream component must attempt to maintain the Link at the Target Link Speed, or at the highest speed supported by both components on the Link (as reported by the values in the training sets – see Section 4.2.4.1), whichever is lower.

~~If Any Upstream Port or Downstream Port with~~ the Hardware Autonomous Speed Disable bit in the Link Control 2 register ~~is clear, the component~~ is permitted to autonomously ~~change-~~
~~adjust~~ the Link speed using implementation specific criteria.

In Section 6.11, page 586, line 20, make the following changes:

Software is permitted to cause a Link to transition to the Polling.Compliance LTSSM state [at a particular speed](#) by writing [the Link Control 2 register in both components with the same value in](#) ~~to~~ the Target Link Speed field and ~~s~~Setting the Enter Compliance bit ~~in the Link Control 2 register in both components~~, and then initiating a Hot Reset on the Link ~~(through the Downstream Port Upstream component)~~. ~~Software is required to write the same value into the Target Link Speed field in both the Upstream and Downstream components.~~

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

In Section 7.8.19, Table 7-28, page 742, Bits 3:0, make the following changes:

Table 7-28: Link Control 2 Register

Bit Location	Register Description	Attributes
3:0	<p>Target Link Speed – For Downstream Ports, this field sets an upper limit on Link operational speed by restricting the values advertised by the Upstream component in its training sequences.</p> <p>...</p> <p>For both Upstream and Downstream Ports, this field is used to set the target compliance mode speed when software is using the Enter Compliance bit to force a Link into compliance mode.</p> <p><u>For Upstream Ports, if the Enter Compliance bit is clear, this field is permitted to have no effect.</u></p> <p>For a Multi-Function device associated with an Upstream Port, the field in Function 0 is of type RWS, and only Function 0 controls the component's Link behavior. In all other Functions of that device, this field is of type RsvdP.</p> <p>Components that support only the 2.5 GT/s speed are permitted to hardwire this field to 0000b.</p>	RWS/RsvdP (see description)

B229 Corrupt Received DLLP

In Section 3.5.2.2, page 200, line 14, make the following changes:

3.5.2.2. Handling of Received DLLPs

...

- A corrupt received DLLP is discarded. This is a Bad DLLP error and, ~~and~~ is a reported error associated with the Port (see Section 6.2).

B230 Loopback

In Section 4.2.2.5, page 238, line 10, make the following changes:

4.2.2.5. Loopback with 128b/130b Code

When using 128b/130b encoding, Loopback Masters must transmit Blocks with the defined 01b and 10b Sync Headers. However, they are not required to transmit an SDS Ordered Set when

transitioning from Ordered Set Blocks to Data Blocks, nor are they required to transmit an EDS Token when transitioning from Data Blocks to Ordered Set Blocks. Masters must transmit SKP Ordered Sets periodically as defined in Section 4.2.7, and they must be capable of processing received (looped-back) SKP Ordered Sets of varying length. Masters are permitted to transmit Electrical Idle Exit Ordered Sets (EIEOS) as defined in Section 4.2.2.2.1. Masters are permitted to transmit any payload in Data Blocks and Ordered Set Blocks that they expect to be looped-back. ~~However, Ordered Set Blocks that match the definition of SKP OS, EIEOS, or EIOS should be avoided since they have defined purposes while in Loopback.~~ If the Loopback Master transmits an Ordered Set Block whose first symbol matches the first symbol of SKP OS, EIEOS, or EIOS, that Ordered Set Block must be a complete and valid SKP OS, EIEOS, or EIOS.

When using 128b/130b encoding, Loopback Slaves must ...

In Section 4.2.6.10.2, page 329, line 24, make the following changes:

4.2.6.10.2. Loopback.Active

- ☐ The loopback master must send valid encoded data. The loopback master must not transmit ~~should avoid sending~~ EIOS as data until it wants to exit Loopback. When operating with 128b/130b encoding, loopback masters must follow the requirements of Section 4.2.2.5.

B231 Flow Control Information Tracked by Receiver

In Section 2.6.1.2, page 161, line 3, make the following changes:

2.6.1.2. FC Information Tracked by Receiver

...

- ☐ When the Link is in the L0 or L0s Link state, Update FCPs for each enabled type of non-infinite FC credit must be scheduled for transmission at least once every 30 μ s (-0%/+50%), except when the Extended Sync bit of the ~~Control~~ Link Control register is set, in which case the limit is 120 μ s (-0%/+50%).

B233 M.2 Specification: Reference to Link Capabilities register

In the M.2 Specification, Revision 1.0, Section 3.1.3.3, page 100, line 5, make the following changes:

3.1.3.3. Clock Request Support Reporting and Enabling

Support for the CLKREQ# dynamic clock protocol should be reported using bit 18 in the PCI Express Link Capabilities register (offset ~~0C4h-00Ch~~). To enable dynamic clock management, bit 8 of the Link Control register (offset 010h) is provided.

B234 WAKE# and OBFF

In the Section 6.19, page 625, line 5, make the following changes:

An OBFF Message received at a Port that does not implement OBFF or when OBFF is not enabled must be handled as an Unsupported Request (UR). [This is a reported error associated with the receiving Port \(see Section 6.2\). If a Port has OBFF enabled using WAKE# signaling, and that Port receives an OBFF Message, the behavior is undefined.](#)

B235 Readiness Time Reporting example times

In Section 7.35, page 912, line 12, make the following changes:

...

Registers and fields in the Readiness Time Reporting Extended Capability are shown in Figure 7-179. Time values are encoded in floating point as shown in Figure 7-180. The actual time value is $Value \times Mul[Scale]$. For example, the value A1Eh represents about 1 second (actually 1.006 sec) and the value ~~40Ah~~ [80Ah](#) represents about 10 ms (actually 10.240 ms).

In Section 7.35.3, page 915, Table 7-158, make the following changes:

Table 7-158: Readiness Time Reporting 2 Register

Bit Location	Register Description	Attributes
23:12	D3_{hot} to D₀D0 Time – If Immediate_Readiness_on_Return_to_D0 is Clear, D3 _{hot} to D0 Time is the time that the Function requires after it is directed from D3 _{hot} to D0 before it is Configuration-Ready and has returned to either D0 _{uninitialized} or D0 _{active} state (see the PCI Bus Power Management Interface Specification). This field is RsvdP if the Immediate_Readiness_on_Return_to_D0 bit is Set. This field is undefined when the Valid bit is Clear. This field must be less than or equal to the encoded value 40Ah 80Ah .	HwInit/RsvdP

B236 *SR-IOV: Changing NumVFs or ARI Capable Hierarchy*

In the SR-IOV Specification, Section 2.1.2, page 27, line 3, make the following changes:

IMPLEMENTATION NOTE:

NumVFs and ARI Capable Hierarchy

After configuring NumVFs and ARI Capable Hierarchy [where applicable](#), software may read First VF Offset and VF Stride to determine how many Bus Numbers would be consumed by the PF's VFs. The additional Bus Numbers, if any, are not actually used until VF Enable is Set

B237 *SR-IOV: Bus Numbering inside a Root Complex*

In the SR-IOV Specification, Section 2.1.2, page 27, line 16, make the following changes:

As in the PCI Express Base Specification, SR-IOV capable Devices [that are associated with an Upstream Port](#) capture the Bus Number from any Type 0 Configuration [Write](#) Request. SR-IOV capable Devices do not capture the Bus Number from any Type 1 Configuration Write Requests. [SR-IOV capable Root Complex Integrated Endpoints use an implementation specific mechanism to assign their Bus Numbers.](#)

B242 *ATS: Terminology: PRG Response PASID Required*

In the PASID ECN to the ATS Specification, Part I, line 15, make the following changes:

In addition, three new bits are defined in the ATS Translation Completion Data Entry (Global Mapping, Execute Permitted, Privileged Mode Access), one new bit is defined in the ATS Invalidation Request Message (Global Invalidate), one new bit is defined in the ATS Extended Capability (Global Invalidate Supported), and one new bit is defined in the Page Request Extended Capability (PRG Response PASID [Required](#)~~Enable~~).

In the PASID ECN to the ATS Specification, in the added Section 4.2.2, page 13, line 29, make the following changes:

*Note: **Green** text was inserted by the ECN and is unchanged by this erratum.*

4.2.2. PASID TLP Prefix Usage

If a Page Request has a PASID TLP Prefix, the corresponding PRG Response Message may optionally contain one as well.

If the PRG Response PASID [Required](#)~~Enable~~ bit is Clear, PRG Response Messages do not have a PASID TLP Prefix.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

If the PRG Response PASID ~~Required~~~~Enable~~ bit is Set, PRG Response Messages have a PASID TLP Prefix if the Page Request also had one. The Function is permitted to use the PASID value from the prefix in conjunction with the PRG Index to match requests and responses.

In PASID TLP Prefixes attached to PRG Response Messages, the Execute Requested and Privileged Mode Requested bits are Reserved and the PASID value is copied from the PASID value of the Page Request.

B246 Device Serial Number

In Section 7.12, page 785, line 8, make the following changes:

It is permitted but not recommended for Root Complex Integrated Endpoints to implement this Capability.

All multi-Function devices that implement this Capability must implement it for Function 0; other Functions that implement this Capability must return the same Device Serial Number value as that reported by Function 0.

A PCI Express ~~multi-device~~ component containing multiple Devices such as a PCI Express Switch that implements this Capability must return the same Device Serial Number for each device.

B248 ~~ATS~~: Page Request Interface and Dirty Bits

In the ATS Specification, Section 1.2, page 19, line 22, make the following changes:

1.2. Page Request Interface Extension

...

The ability to page, begs the question of page table status flag management. Typical TAs associate flags (e.g. dirty and access indications) with each untranslated address. Without any additional hints about how to manage pages mapped to an ~~I/O device Function~~, ~~an RC~~such TAs would need to conservatively assume that ~~if an I/O device can write to a page, it has written to the page when they grant a Function permission to read or write a page, that Function will use the permission.~~ ~~I/O~~ Such writable pages would need to be marked as dirty before their translated addresses are made available to a ~~device~~Function.

This conservative dirty-on-write-permission-grant behavior is generally not a significant issue for Functions that do not support paging ~~non-pageable devices~~, where ~~I/O~~ pages are pinned and the cost of saving a clean page to memory will seldom be paid. However, ~~devices~~ Functions that support the Page Request Interface could pay a significant penalty if all writable pages are treated as dirty, since such ~~devices~~ Functions operate without pinning their accessible memory footprints and may issue speculative page requests for performance. The cost of saving clean pages (instead of just discarding them) in such systems can diminish the value of otherwise attractive paging techniques. This can cause significant performance issues and risk functional issues in circumstances where the backing store is unable to be written, such as a CD-ROM.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

The No Write (NW) flag in Translation Requests indicates ~~that a Function whether a page should be marked as dirty if it is writable, or whether a device~~ is willing to restrict its usage to only reading the page, independent of the access rights that would otherwise have been granted.

If a device chooses to request only read access ~~by issuing a Translation Request with the NW flag Set and later then~~ determines that it ~~does wish~~ needs to write to the page, then the device ~~must~~ ~~would need to~~ issue a new Translation Request.

Upon receiving a Translation Request with the NW flag Clear, TAs are permitted to mark the associated pages dirty. It is strongly recommended that Functions not issue such Requests unless they have been given explicit write permission. An example of write permission is where the host issues a command to a Function to load data from a storage device and write that data into memory.

...

In the ATS Specification, Section 2.2.5, page 24, line 23, make the following changes:

2.2.5. No Write (NW) Flag

The No Write flag, when Set, indicates that the Function is requesting read-only access for this translation.

The TA may ignore the No Write Flag, however, if the TA responds with a translation marked as read-only then the Function must not issue MemWrite transactions using that translation. In this case, the Function may issue another translation request with the No Write flag Clear, which may result in a new translation completion with or without the W (Write) bit Set.

Upon receiving a Translation Request with the NW flag Clear, TAs are permitted to mark the associated pages dirty. It is strongly recommended that Functions not issue such Requests unless they have been given explicit write permission.

In the ATS Specification, Section 4.1, Table 4-1, page 41, make the following changes:

4.1. Page Request Message

...

Table 4-1: Page Request Message Data Fields

Field	Meaning
W	<p>Write Access Requested – This field, when Set, indicates that the requesting Function seeks write access and/or zero-length read access to the associated page. When Clear, this field indicates that the requesting Function will not write to the associated page.</p> <p><u>Upon receiving a Page Request Message with the W field Set, the host is permitted to mark the associated page dirty. Thus, Functions must not issue such Requests unless the Function has been given explicit write permission.</u></p>

B249 Access Control Services for SR-IOV

In Section 6.12, page 586, line 30, make the following changes:

6.12. Access Control Services (ACS)

ACS defines a set of control points within a PCI Express topology to determine whether a TLP ~~should~~ **is to** be routed normally, blocked, or redirected. ACS is applicable to RCs, Switches, and multi-Function devices.⁸⁹ For ACS requirements, single-Function devices that are SR-IOV capable must be handled as if they were multi-Function devices, since they essentially behave as multi-Function devices after their Virtual Functions (VFs) are enabled.

In Section 6.12.1.2, page 590, line 17, make the following changes:

6.12.1.2. ACS Functions in SR-IOV Capable and Multi-Function Devices

This section applies to multi-Function device ACS Functions, with the exception of Downstream Port Functions, which are covered in the preceding section. For ACS requirements, single-Function devices that are SR-IOV capable must be handled as if they were multi-Function devices.

- ☐ ACS Source Validation: must not be implemented.
- ☐ ACS Translation Blocking: must not be implemented.
- ☐ ACS P2P Request Redirect: must be implemented by Functions that support peer-to-peer traffic with other Functions. This includes SR-IOV Virtual Functions (VFs).

ACS P2P Request Redirect is subject to interaction with the ACS P2P Egress Control and ACS Direct Translated P2P mechanisms (if implemented). Refer to Section 6.12.3 for more information.

When ACS P2P Request Redirect is enabled in a multi-Function or SR-IOV capable device, peer-to-peer Requests (between Functions of the device) must be redirected Upstream towards the RC.

Completions are never affected by ACS P2P Request Redirect.

- ☐ ACS P2P Completion Redirect: must be implemented by Functions that implement ACS P2P Request Redirect.

The intent of ACS P2P Completion Redirect is to avoid ordering rule violations between Completions and Requests when Requests are redirected. Refer to Section 6.12.5 for more information.

ACS P2P Completion Redirect does not interact with ACS controls that govern Requests.

When ACS P2P Completion Redirect is enabled in a multi-Function or SR-IOV capable device, peer-to-peer Read Completions that do not have the Relaxed Ordering bit set must be redirected Upstream towards the RC. Otherwise, peer-to-peer Completions must be routed normally.

Requests are never affected by ACS P2P Completion Redirect.

- ☐ ACS Upstream Forwarding: must not be implemented.

- ❑ ACS P2P Egress Control: implementation is optional; is based on Function Numbers or Function Group Numbers; controls peer-to-peer Requests between the different Functions within the multi-Function [or SR-IOV capable](#) device.

ACS P2P Egress Control is subject to interaction with the ACS P2P Request Redirect and ACS Direct Translated P2P mechanisms (if implemented). Refer to Section 6.12.3 for more information.

Each Function within a multi-Function [or SR-IOV capable](#) device that supports ACS P2P Egress Control can be selectively enabled to block peer-to-peer communication with other Functions or Function Groups⁹⁹ within the device. This is configured on a per Function basis.

With ACS P2P Egress Control in multi-Function [or SR-IOV capable](#) devices, controls in the “sending” Function determine if the Request is blocked, and if so, the “sending” Function handles the ACS Violation error per Section 6.12.4.

When ACS Function Groups are enabled in an ARI Device, ACS P2P Egress Controls are enforced on a per Function Group basis instead of a per Function basis. See Section 6.13.

Completions are never affected by ACS P2P Egress Control.

- ❑ ACS Direct Translated P2P: must be implemented if the multi-Function [or SR-IOV capable](#) device Function supports Address Translation Services (ATS) and also peer-to-peer traffic with other Functions.

When ACS Direct Translated P2P is enabled in a multi-Function [or SR-IOV capable](#) device Function, peer-to-peer Memory Requests whose Address Type (AT) field indicates a Translated address must be routed normally (“directly”) to the peer Function, regardless of ACS P2P Request Redirect and ACS P2P Egress Control settings. All other peer-to-peer Requests must still be subject to ACS P2P Request Redirect and ACS P2P Egress Control settings.

Completions are never affected by ACS Direct Translated P2P.

In Section 6.12.1.3, page 591, line 31, make the following changes:

6.12.1.3. Functions in Single-Function Devices

This section applies to single-Function device Functions, with the exception of Downstream Port Functions [and SR-IOV capable Functions](#), which are covered in ~~a~~-preceding sections. [For ACS requirements, single-Function devices that are SR-IOV capable must be handled as if they were multi-Function devices.](#)

No ACS capabilities are applicable, and the Function must not implement an ACS Extended Capability structure.

B250 *Enhanced Allocation: Address Maps*

In the Enhanced Allocation (EA) ECN, Section 6.2.5.1, page 4, line 8, make the following changes:

*Note: **Green** text was inserted by the ECN and is unchanged by this erratum.*

6.2.5.1. Address Maps

After determining this information, power-up software can map the I/O controllers into reasonable locations and proceed with system boot. In order to do this mapping in a device independent manner, the base registers for this mapping are placed in the predefined header portion of Configuration Space. **It is strongly recommended that power-up firmware/software also support the optional Enhanced Allocation mechanism (see Section 6.9).** Resources allocated by a specific Function using Enhanced Allocation must be handled by system firmware/software as being unavailable for use by any other entity, regardless of the state of that specific Function, including cases where that Function is disabled, and without regard to the state of that Function's Memory Space Enable, IO Space Enable bits in the Command register, or any function-specific mechanisms.

...

Devices that map control functions into I/O Space must not consume more than 256 bytes per I/O Base Address register **or per each I/O Space entry in the Enhanced Allocation capability.**

...

B252 *Latency Tolerance Reporting (LTR)*

In Section 6.18, page 619, line 15, make the following changes:

6.18. Latency Tolerance Reporting (LTR) Mechanism

...

When the LTR Mechanism Enable bit is cleared, if a device's most recently sent LTR Message (since the last DL_DOWN to DL_Up transition) reported latency tolerance values with any Requirement bit set, then ~~it~~ **one of the following applies.**

- ☐ If the bit was cleared due to a Configuration Write to the Device Control 2 register, the device must send a new LTR Message with all the Requirement bits clear.
- ☐ If the bit was cleared due to an FLR, it is strongly recommended that the device send a new LTR Message with all the Requirement bits clear.

When a Downstream Port goes to DL_Down status, any previous latencies recorded for that Port must be treated as invalid.

An LTR Message from a device ...

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

In Section 6.18, page 621, line 3, make the following changes:

- ❑ ~~When a Switch Downstream Port goes to DL_Down status, the previous latencies recorded for that Port must be treated as invalid.~~ A new LTR message must be transmitted Upstream if the conglomerated latencies are changed as a result of DL_Down invalidating the previous latencies recorded for that Port.

B253 LCRC and Sequence Number Rules (TLP Transmitter)

In Section 3.5.2.1, page 196, line 26, make the following changes:

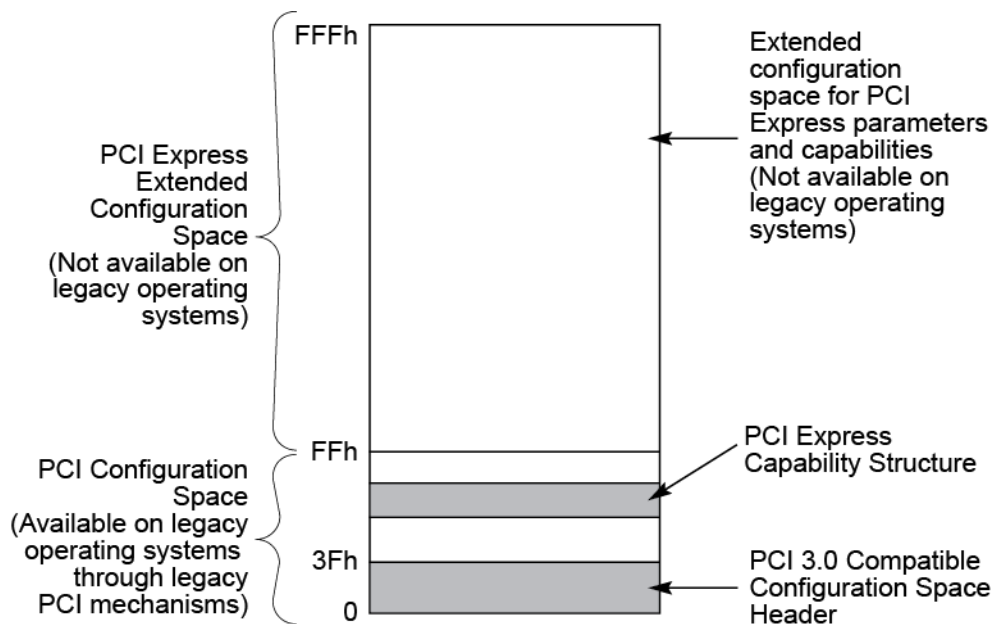
AckFactor represents the number of maximum size TLPs which can be received before an Ack is sent, and is used to balance Link bandwidth efficiency and retry buffer size – the value varies according to Max_Payload_Size and Link width, and is defined in Table 3-7, Table 3-8, and Table ~~3-6~~.

B254 PCI Express Capability Structure

In Section 3.5.2.1, page 655, Figure 7-3, change text as following:

PCI Express Capability Structure ~~Capability needed by BIOS or by driver software on non-PCI Express aware operating systems~~

so the figure becomes the following:



OM14301B

Figure 7-3: PCI Express Configuration Space Layout

B255 *Enhanced Allocation: Reserved Encodings*

In the Enhanced Allocation (EA) ECN, page 13, Table 6-1, make the following changes:

Value (h)	Resource Definition
...	...
08-FCE	Reserved for future use; System firmware/software must not write to this entry, and must not attempt to interpret this entry or to use this resource. When software reads a Primary Properties value that is within this range, is it strongly recommended that software treat this resource according to the value in the Secondary Properties

B256 *Enhanced Allocation: Virtual Function BEI Ranges*

In the Enhanced Allocation (EA) ECN, page 11, line 14, make the following changes:

- For Memory ~~or I/O~~ BARs where the Primary or Secondary ~~Property~~ Properties is 03h or 04h it is permitted to assign the same BEI in the range of 0-5 once for a range where Base + MaxOffset is below 4GB, and again for a range where Base + MaxOffset is greater than 4GB; It is not otherwise permitted to assign the same BEI in the range 0-5 for more than one entry.
- For Virtual Function BARs where the Primary or Secondary ~~Property~~ Properties is 03h or 04h it is permitted to assign the same BEI in the range of ~~0-5~~ 9-14 once for a range where Base + MaxOffset is below 4GB, and again for a range where Base + MaxOffset is greater than 4GB; It is not otherwise permitted to assign the same BEI in the range ~~0-5-9-14~~ for more than one VF entry.

B257 *Lane Equalization Control Register Nomenclature*

In Section 6.6.2, page 562, line 31, make the following changes:

- ~~Link Lane~~ Equalization Control register in the Secondary PCI Express Extended Capability structure

In Section 7.27, page 857, Figure 7-132, change the name of the field at offset 0Ch as follows:

Lane Equalization Control Register (Sized by Maximum Link Width)

so that the figure looks like:

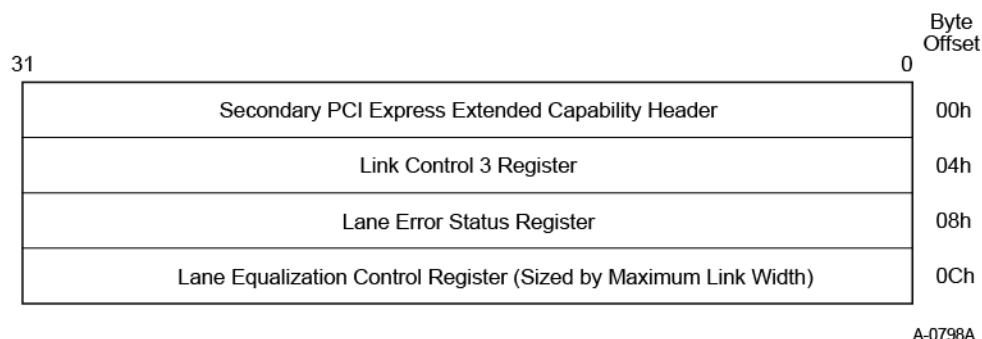


Figure 7-132: Secondary PCI Express Extended Capability Structure

In Section 7.27.4, page 860, line 4, make the following changes:

The [Lane](#) Equalization Control register consists of control fields required for per Lane equalization and the number of entries in this register are sized by Maximum Link Width (see Section 7.8.6). Each entry contains the values for the Lane with the corresponding default Lane number which is invariant to Link width and Lane reversal negotiation that occurs during Link training.

In Section 7.27.4, page 861, change the caption for Figure 7-137 as follows:

Figure 7-137: Lane ~~((Maximum Link Width—1):0)~~ Equalization Control Register [Entry](#)

In Section 7.27.4, page 861, change the caption for Table 7-117 as follows:

Table 7-117: Lane ~~((Maximum Link Width—1):0)~~ Equalization Control Register [Entry](#)

In Section 4.2.3, page 241, line 1, make the following changes:

rate. These preset values are derived from the Upstream Port Transmitter Preset and Upstream Port Receiver Preset Hint fields of each Lane^{2s} Equalization Control register [entry](#). After the data rate change to 8.0 GT/s, the Upstream Port transmits TS1 Ordered Sets with the preset values it received. The preset values must be within the operable range defined in Section 4.3.3.5.2 if reduced swing will be used by the Transmitter.

Phase 1: Both components make the Link operational enough at 8.0 GT/s data rate to be able to exchange TS1 Ordered Sets to complete remaining phases for the fine-tuning their Transmitter/Receiver pairs. It is expected that the Link will operate at a BER of less than 10^{-4} before the component is ready to move on to the next Phase.

The Downstream Port initiates Phase 1 by transmitting TS1 Ordered Sets with EC=01b (indicating Phase 1) to the Upstream Port using the preset values in the Downstream Port Transmitter Preset and, optionally, the Downstream Port Receiver Preset Hint fields of each Lane^{2s} Equalization Control register [entry](#). The Upstream Port, after ...

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

In Section 4.2.6.4.4, page 310, line 1, make the following changes:

The Downstream Port must transmit EQ TS2 Ordered Sets (TS2 Ordered Sets with Symbol 6 bit 7 set to 1b) with the Transmitter Preset and Receiver Preset Hint fields set to the values specified by the Upstream Port Transmitter Preset and the Upstream Port Receiver Preset Hint fields of each Lane's Equalization Control register [entry](#) if all of the following conditions are satisfied: ...

B258 Timing Recommendations for Latency Tolerance Reporting

In section 7.25.2, page 851, Table 7-108, make the following changes:

Bit Location	Register Description	Attributes
9:0	Max Snoop LatencyValue – Along with the Max Snoop LatencyScale field, this register specifies the maximum snoop latency that a device is permitted to request. Software should set this to the platform's maximum supported latency or less. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms. The default value for this field is 00 0000 0000b.	RW
12:10	Max Snoop LatencyScale – This register provides a scale for the value contained within the Maximum Snoop LatencyValue field. Encoding is the same as the LatencyScale fields in the LTR Message. See Section 6.18. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms. The default value for this field is 000b. Hardware operation is undefined if software writes a Not Permitted value to this field.	RW

In section 7.25.3, page 852, Table 7-109, make the following changes:

Bit Location	Register Description	Attributes
9:0	Max No-Snoop LatencyValue – Along with the Max No-Snoop LatencyScale field, this register specifies the maximum no-snoop latency that a device is permitted to request. Software should set this to the platform's maximum supported latency or less. It is strongly	RW

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

	recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms. The default value for this field is 00 0000 0000b.	
12:10	Max No-Snoop LatencyScale – This register provides a scale for the value contained within the Max No-Snoop LatencyValue field. Encoding is the same as the LatencyScale fields in the LTR Message. See Section 6.18. It is strongly recommended that any updates to this field are reflected in LTR Message(s) sent by the device within 1 ms. The default value for this field is 000b. Hardware operation is undefined if software writes a Not Permitted value to this field.	RW

In Section 6.18, page 623, line 13, change the implementation note as follows:

IMPLEMENTATION NOTE:

Optimal Use of LTR

...

Note that the RC may delay processing of device Request TLPs, provided it satisfies the device's service requirements. If, for example, an Endpoint connected to Root Port 1 reports a latency tolerance of 100 μ s, and an Endpoint on Root Port 2 report a value of 30 μ s, the RC might implement a policy of stalling an initial Request following an idle period from Root Port 1 for 70 μ s before servicing the Request with a 30 μ s latency, thus providing a perceived service latency to the first Endpoint of 100 μ s. This RC behavior provides the RC the ability to batch together Requests for more efficient servicing.

[It is possible that, after it is determined that the RC can service snoop and no-snoop Requests from all Endpoints within the maximum snoop and maximum no-snoop time intervals, this information may be communicated to Endpoints by updating the Max Snoop LatencyValue, Max Snoop LatencyScale and Max No-Snoop LatencyValue, Max No-Snoop LatencyScale fields. The intention of this communication would be to prevent Endpoints from sending needless LTR updates.](#)

[When an Endpoint's LTR value for snoop Requests changes to become larger \(looser\) than the value indicated in the Max Snoop LatencyValue/Scale fields, it is recommended that the Endpoint send an LTR message with the snoop LTR value indicated in the Max Snoop LatencyValue/Scale fields. Likewise, when an Endpoint's LTR value for no-snoop Requests changes to become larger \(looser\) than the value indicated in the Max No-Snoop LatencyValue/Scale fields, it is recommended that the Endpoint send an LTR message with the no-snoop LTR value indicated in the Max No-Snoop LatencyValue/Scale fields.](#)

It is recommended that Endpoints buffer Requests as much as possible, and then use the full Link bandwidth in bursts that are as long as the Endpoint can practically support, as this will generally lead to the best overall platform power efficiency.

B259 Completion Timeout Mechanism

In Section 2.8, page 170, line 26, make the following changes:

The Completion Timeout mechanism may be disabled by configuration software. The Completion Timeout limit is set in the Completion Timeout Value field of the Device Control 2 register. ~~Refer to Section 2.2.8.10.~~ A Completion Timeout is a reported error associated with the Requester Function (see Section 6.2).

B260 Slot Power Limit

In Section 7.15.3, page 804, Table 7-65, make the following changes:

Bit Location	Register Description	Attributes
7:0	<p>Base Power – Specifies in watts the base power value in the given operating condition. This value must be multiplied by the data scale to produce the actual power consumption value except when the Data Scale field equals 00b (1.0x) and Base Power exceeds EFh, the following alternative encodings are used:</p> <p>F0h = greater than 239 W and less than or equal to 250 W Slot Power Limit</p> <p>F1h = greater than 250 W and less than or equal to 275 W Slot Power Limit</p> <p>F2h = greater than 275 W and less than or equal to 300 W Slot Power Limit</p> <p>F3h to FFh = Reserved for Slot Power Limit values above greater than 300 W</p>	RO

B263 AtomicOp Completion Rules

In Section 6.15.2, page 610, line 20, make the following changes:

6.15.2 AtomicOp Transaction Protocol Summary

...

- ❑ Unless there's a higher precedence error, an [AtomicOp-aware](#) Completer must handle a Poisoned AtomicOp Request as a Poisoned TLP Received error, and must also return a Completion with a Completion Status of Unsupported Request (UR). See Section 2.7.2.2. The value of the target location must remain unchanged.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

- ❑ If the Completer of an AtomicOp Request encounters an uncorrectable error accessing the target location or carrying out the Atomic operation, the Completer must handle it as a Completer Abort (CA). The subsequent state of the target location is implementation specific.
- ❑ [AtomicOp-aware](#) Completers are required to handle any properly formed AtomicOp Requests with types or operand sizes they don't support as an Unsupported Request (UR). If the Length field in an AtomicOp Request contains an unarchitected value, the Request must be handled [by an AtomicOp-aware Completer](#) as a Malformed TLP. See Section 2.2.7.
- ❑ If any Function in a multi-Function device supports AtomicOp Completer or AtomicOp routing capability, all Functions with Memory Space BARs in that device must decode properly formed AtomicOp Requests and handle any they don't support as an Unsupported Request (UR). Note that in such devices, Functions lacking AtomicOp Completer capability are forbidden to handle properly formed AtomicOp Requests as Malformed TLPs.
- ❑ If an RC has any Root Ports that support AtomicOp routing capability, all Root Complex Integrated Endpoints in the RC reachable by forwarded AtomicOp Requests must decode properly formed AtomicOp Requests and handle any they don't support as an Unsupported Request (UR).
- ❑ With an AtomicOp Request having a supported type and operand size, the [AtomicOp-aware](#) Completer is required either to carry out the Request or handle it as Completer Abort (CA) for any location in its target Memory Space. Completers are permitted to support AtomicOp Requests on a subset of their target Memory Space as needed by their programming model (see Section 2.3.1). Memory Space structures defined or inherited by PCI Express (e.g., the MSI-X Table structure) are not required to be supported as AtomicOp targets unless explicitly stated in the description of the structure.

B264 PASID Completions

In Section 6.20.2.1, page 630, line 15, make the following changes:

For Endpoints, the following rules apply:

- ❑ The Endpoint is not permitted to send TLPs with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$.
- ❑ The Endpoint is optionally permitted to signal an error when it receives a ~~TLPRequest~~ with a PASID value greater than or equal to $2^{\text{Max PASID Width}}$. ~~For Requests, this error~~ This is an Unsupported Request [error associated with the Receiving Port \(see Section 6.2\)](#) ~~and for Completions, this error is Unexpected Completion.~~

For Root Complexes, the following rules apply:

- ❑ A Root Complex is not permitted to send a TLP with a PASID value greater than it supports.
- ❑ A Root Complex is optionally permitted to signal an error when it receives a ~~TLPRequest~~ with a PASID value greater than it supports. ~~For Requests, this error~~ This is an Unsupported Request [error associated with the Receiving Port \(see Section 6.2\)](#) ~~and for Completions, this error is Unexpected Completion.~~

B265 *SR-IOV: ARI Capable Hierarchy*

In the SR-IOV Specification, page 39 bit 4, make the following changes:

PCI Express Endpoint:

... RW in the lowest-numbered PF of the Device and hardwired to 0b in all other PFs.

[This bit is not affected by FLR of any PF or VF.](#)

...

In the SR-IOV Specification, page 42 line 9, make the following changes:

A Device may use the setting of ARI Capable Hierarchy to determine the values for First VF Offset (see Section 3.3.9) and VF Stride (see Section 3.3.10). The effect of changing ARI Capable Hierarchy is undefined if VF Enable is Set in any PF. [This bit must be set to its default value upon Conventional Reset.](#) This bit is not affected by FLR of any PF or VF.

C266 *LNR Control Register Offset*

In Section 7.30.3, change the section header as follows:

7.30.3. LNR Control Register (Offset **04h06h)**

C267 *Resizable BAR Control Register Name Usage*

In Section 7.22, Page 839, line 12 make the following changes:

IMPLEMENTATION NOTE

Using the Capability During Resource Allocation

...

Software then writes the size to the Resizable BAR [Control](#) ~~Command~~ register for the appropriate BAR for the Function.

C268 *Resizable BAR Capabilities Register Name Usage*

In Section 4.2.6.6.1, Page 318, line 22, make the following changes:

A Receiver must implement L0s if its Port advertises support for L0s, as indicated by the ASPM Support field in the Link [Capabilities](#) ~~Capability~~ register.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

In Section 4.2.6.6.2, Page 320, line 5, make the following changes:

A Transmitter must implement L0s if its Port advertises support for L0s, as indicated by the ASPM Support field in the Link [Capabilities](#) ~~Capability~~ register.

In Section 5.4.1.1, Page 467, line 6, make the following changes:

Such software might not even check the ASPM Support field in the Link [Capabilities](#) ~~Capability~~ register, might not recognize ...

B269 DC Balance Tracking

In Section 4.2.4.1, Page 247, line 11, make the following changes:

When using 128b/130b encoding, Transmitters are required to track the running DC Balance of the bits transmitted on the wire (after scrambling) ~~for~~ [that constitute the](#) TS1 and TS2 Ordered Sets [only](#).

A270 ECR and ECN Process

In the PCISIG ECR and ECN Process Overview document, page 2 line 3 make the following changes:

At any one time, there can be only one current version of all PCI specifications and design guides, termed the “production version.” [ECNs and errata become part of an associated specification at the moment they are released as final into the Members Specification area](#). At the same time, there may also be proposed revisions to the specifications and design guides that are not yet approved and are held confidential (as deemed necessary by the Board of Directors) within the PCI-SIG Workgroups.

A271 Expanded Resizable BAR: BAR Size field

In the Expanded Resizable BARs ECN, page 6, bits 13:8 row, make the following changes:

BAR Size – This ...

When this register field is programmed, the value is immediately reflected in the size of the resource, as encoded in the number of read-only bits in the BAR.

Software must only write values that correspond to those indicated as supported in the Resizable BAR Capability and Control registers. Writing an unsupported value will produce undefined results. [BAR Size bits that never need to be Set in order to indicate every supported size are permitted to be hardwired to 0.](#)

B274 Advisory Non-Fatal

In Section 6.2.7, Table 6-2, page 525, ECRC Check Failed row, make the following changes:

Error Name	Error Type (Default Severity)	Detecting Agent Action	References
ECRC Check Failed	Uncorrectable (Non-Fatal)	Receiver (if ECRC checking is supported): Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Section Sections 6.2.3.2.4.1 and 6.2.3.2.4.2 . Log the prefix/header of the TLP that encountered the ECRC error.	Section 2.7.1

In Section 6.2.7, Table 6-2, page 526, ACS Violation row, make the following changes:

Error Name	Error Type (Default Severity)	Detecting Agent Action	References
ACS Violation	Uncorrectable (Non-Fatal)	Receiver (if checking): Send ERR_NONFATAL to Root Complex or ERR_COR for the Advisory Non-Fatal Error case described in Sections 6.2.3.2.4.1 . Log the prefix/header of the Request TLP that encountered the error.	

C275 MFVC term usage

In Section 2.5.1, page 149, line 38 make the following changes:

2.5.1 Virtual Channel Identification (VC ID)

If a multi-Function device implements an MFVC Capability structure, its VC hardware resources are distinct from the VC hardware resources associated with any VC Capability structures of its Functions. The VC ID uniqueness requirement (first bullet above) still applies individually for the ~~MVFC~~[MFVC](#) and any VC Capability structures.

B276 DRS and FRS Message Subtype

In Section 2.2.8.6.3, page 108, line 14 make the following changes:

2.2.8.6.3. Device Readiness Status (DRS) Message

...

- ☐ The TLP Type must be Msg.
- ☐ The TC[2:0] field must be 000b.
- ☐ The Attr[2:0] field is Reserved.
- ☐ The Tag field is Reserved.
- ☐ [The Subtype field must be 08h.](#)
- ☐ The Message Routing field must be set to 100b – Local – Terminate at Receiver.

In Section 2.2.8.6.4, page 109, line 16, make the following changes:

2.2.8.6.4. Function Readiness Status (FRS) Message

...

- ☐ The TLP Type must be Msg.
- ☐ The TC[2:0] field must be 000b.
- ☐ The Attr[2:0] field is Reserved.
- ☐ The Tag field is Reserved.
- ☐ [The Subtype field must be 09h.](#)
- ☐ The FRS Reason[3:0] field indicates why the FRS Message was generated:

In Table F-2, Page 1031, make the following changes:

Table F-2 PCI-SIG Defined VDM subtype Usage

Subtype	Routing r[2:0]	Type	Description
0000 0000	010 or 011	MsgD	LN Message, see Section 2.2.8.6.1.1
0111 1111 0000 1000	100	Msg	Device Readiness Status, see Section 2.2.8.6.3
0111 1111 0000 1001	000	Msg	Function Readiness Status, see Section 2.2.8.6.4

B277 **ATS:** Translation Completion

In the ATS Specification, Section 2.3, page 25 Figure 2-5 make the following changes:




Figure 2-5: Translation Completion with No Data



IMPLEMENTATION NOTE

Byte Count Field for Unsuccessful Translation Completions with No Data

Previous versions of this specification indicated the Byte Count and Lower Address field should be 0000 0000 0000b for Unsuccessful Translation Completions with No Data. It is strongly recommended that implementations do not depend on the Byte Count and Lower Address field being set to any particular value in Unsuccessful Translation Completions with No Data.

The values and meaning for the Completion Status field are listed in Table 2-2.

B278 **ATS:** Translation Completion Length

In the ATS Specification, Section 2.3, page 26, line 7, make the following changes:

2.3. Translation Completion

... If the results are returned in multiple packets, the first packet will have a Lower Address field of RCB minus ([Total Completion](#) Length * 4) and subsequent packets will have a Lower Address field of 00000000b.

In the ATS Specification, Section 2.4, page 30, line 3, make the following changes:

2.4. Completions with Multiple Translations

An ATC is allowed to request that the TA provide translations for a virtually contiguous range of addresses. It does this by setting the Length field in the Translation Request to a value that is two times the number of requested translations as long as the request size ([Total Completion](#) Length * 4) is not larger than either Max_Read_Request_Size or RCB.

B279 *ATS: Completions with Multiple Translations*

In the ATS Specification, Section 2.4, page 30, line 16, make the following changes:

2.4. Completions with Multiple Translations

...

A [successful](#) Translation Completion ~~may require~~ [must consist of](#) one or two CplDs.

B280 *ATS: Translation Completion Lower Address Field*

In the ATS Specification, Section 2.3, page 26, line 8, make the following changes:

The Lower Address field will contain a value that will make the packet consistent with RCB semantics. If the result is returned in a single packet, Lower Address is set to RCB minus Byte Count. If the results are returned in multiple packets, the first packet will have a Lower Address field of RCB minus ([Total Completion](#) Length * 4) and subsequent packets will have a Lower Address field of 000 0000b. [See section 2.4 for additional requirements for multiple packet completions.](#)

If the Completion ...

In the ATS Specification, Section 2.4, page 30, line 20, make the following changes:

For such a CplD, if the sum of Byte Count and Lower Address is not a multiple of RCB, then the CplD is the last of a sequence. ~~and it is an error if~~ [If](#) no previous CplD [for this request](#) has been received, [an error has occurred](#) and all translation values should be discarded.

B281 *Completer breaking up Responses*

In Section 2.3.1.1, page 133, line 18, make the following changes:

2.3.1.1. Data Return for Read Requests

...

Read Completion Boundary (RCB) determines the naturally aligned address boundaries on which [a Completer is permitted to break up the response for a single Read Request into](#) ~~a Read Request may be serviced with~~ multiple Completions

In Section 2.3.1.1, page 134, line 1, make the following changes:

- ☐ Requests which do cross the address boundaries at integer multiples of RCB bytes [are permitted to](#) ~~may~~ be completed using more than one Completion [subject to the following rules:](#) ~~but the data must not be fragmented except along the following address boundaries:~~











**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

- The first Completion must start with the address specified in the Request, and if successful must end at one of the following:
 - the address that satisfies ~~specified in the Request plus the length specified by the Request (i.e., the entire Request)~~
 - an address boundary between the start and end of the Request at an integer multiple of RCB bytes
- ~~The~~ If the final Completion is successful, it must end ~~with~~ at the address that satisfies the entire ~~specified in the Request plus the length specified by the Request~~

In Section 2.3.1.1, page 134, line 19, make the following changes:

- ❑ Multiple Memory Read Completions for a single Read Request must return data in increasing address order.
- ❑ If all the Memory Read Completions for a single Read Request have a Successful Completion Status, the sum of their payloads must equal the size requested.
- ❑ For each Memory Read Completion, the Byte Count field must indicate the remaining number of bytes required to complete the Request including the number of bytes returned with the Completion, except when the BCM bit is 1b.²¹
 - The total number of bytes required to complete a Memory Read Request is calculated as shown in Table 2-37.
 - If a Memory Read Request is completed using multiple Completions, the Byte Count value for each successive Completion is the value indicated by the preceding Completion minus the number of bytes returned with the preceding Completion.
- ❑ The Completion Data area begins at the DW address specified by the Request. In the first or only Data DW of the first or only Completion, only the bytes configured as active in the First BE field in the Request contain valid data. Bytes configured as inactive in the First BE field in the Request will return undefined content.
- ❑ In the last Data DW of the last successful Completion, only the bytes configured as active in the Last BE field in the Request contain valid data. Bytes configured as inactive in the Last BE field in the Request will return undefined content.
- ❑ All the Completion Data bytes, including those with undefined content, are included in all CRC calculations.
- ❑ Figure 2.x presents an example of the above. The example assumes a single Completion TLP is returned:

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

Completion Data					Request Byte Enables
Request Address (DW)	Byte 0	Byte 1	Byte 2	Byte 3	
START	undefined content	undefined content	undefined content		First BE: 1000
START + 1					
START + 2					
START + 3		undefined content	undefined content	undefined content	Last BE: 0001

Length = 4d;
Byte Count = 10d;

Figure 2.x Completion Data when some Byte Enable bits are Clear

Change footnote 21 as follows:

²¹ Only PCI-X completers set the BCM bit to 1b. [PCI Express completers are not permitted to set the BCM bit.](#)

B282 Reserved bits in SKP OS Generation Vector

In Section 7.27.2, Table 7-115, page 859, make the following changes:

9:15	<p>Enable Lower SKP OS Generation Vector – When the Link is in L0 and the bit in this field corresponding to the current Link speed is Set, SKP Ordered Sets are scheduled at the rate defined for SRNS, overriding the rate required based on the clock tolerance architecture. See Section 4.2.7 for additional requirements.</p> <p>Bit definitions within this field are:</p> <p>Bit 0 2.5 GT/s Bit 1 5.0 GT/s Bit 2 8.0 GT/s Bits 6:3 RsvdP</p> <p>Bits <u>Each unreserved bit</u> in this field are <u>must be</u> RW if the corresponding bit in the Lower SKP OS Generation Supported Speeds Vector is Set; otherwise, <u>the bit must</u> they are permitted to be <u>RW or</u> hardwired to 0.</p> <p>The default value of this field is 000 0000b.</p>	RW/RsvdP
------	---	----------

B283 *Enhanced Allocation ECN: Enhanced Allocation Capability*

In the Enhanced Allocation ECN, Section 6.9.1.1, page 7, Figure 6-2,

- Change “Reserved” to “RsvdP”, and
- Change “Next Pointer” to “Next Capability Pointer”.

In the Enhanced Allocation ECN, Section 6.9.1.1, page 7, unlabeled table after Figure 6-2, make the following changes:

Bits	Field	Description
7:0	CAP_ID <u>Capability ID</u>	Must be set to 14h to indicate Enhanced Allocation capability. This field is read only.
15:8	NXT_PTR <u>Next Capability Pointer</u>	Pointer to the next item in the capabilities list. Must be NULL for the final item in the list. This field is read only.
21:16	Num Entries	Number of entries following the first DW of the capability. Value of 00 0000b is permitted and means there are no entries. This field is read only.

In the Enhanced Allocation ECN, Section 6.9.1.2, page 8, Figure 6-3, Change “Reserved” to “RsvdP”.

In the Enhanced Allocation ECN, Section 6.9.1.3, page 9, Figure 6-4

- Change “Reserved” to “RsvdP” (Bits 29:24)
- Indicate Bit 3 is “RsvdP”

In the Enhanced Allocation ECN, Section 6.9.1.3, page 10, unlabeled table following Figure 6-4, make the following changes:

Bits	Field	Description
...
7:4	BEI	BAR Equivalent Indicator – This field indicates the equivalent BAR for this entry.

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

		<p>Specific rules for use of this field are given in the text following this table.</p> <table><tr><th>BEI Value</th><th>Description</th></tr><tr><td>0</td><td>Entry is equivalent to BAR at location 10h</td></tr><tr><td>1</td><td>Entry is equivalent to BAR at location 14h</td></tr><tr><td>2</td><td>Entry is equivalent to BAR at location 18h</td></tr><tr><td>3</td><td>Entry is equivalent to BAR at location 1Ch</td></tr><tr><td>4</td><td>Entry is equivalent to BAR at location 20h</td></tr><tr><td>5</td><td>Entry is equivalent to BAR at location 24h</td></tr><tr><td>6</td><td>Permitted to be used by a Function with a Type 1 Configuration Space header function only, optionally used to indicate a resource that is located behind the Function</td></tr><tr><td>7</td><td>Equivalent Not Indicated</td></tr><tr><td>8</td><td>Expansion ROM Base Address</td></tr><tr><td>9-14</td><td>Entry relates to VF BARs 0-5 respectively</td></tr><tr><td>15</td><td>Reserved – Software must treat values in this range as “Equivalent Not Indicated”</td></tr></table> <p>This field is HwInit.</p>	BEI Value	Description	0	Entry is equivalent to BAR at location 10h	1	Entry is equivalent to BAR at location 14h	2	Entry is equivalent to BAR at location 18h	3	Entry is equivalent to BAR at location 1Ch	4	Entry is equivalent to BAR at location 20h	5	Entry is equivalent to BAR at location 24h	6	Permitted to be used by a Function with a Type 1 Configuration Space header function only, optionally used to indicate a resource that is located behind the Function	7	Equivalent Not Indicated	8	Expansion ROM Base Address	9-14	Entry relates to VF BARs 0-5 respectively	15	Reserved – Software must treat values in this range as “Equivalent Not Indicated”
BEI Value	Description																									
0	Entry is equivalent to BAR at location 10h																									
1	Entry is equivalent to BAR at location 14h																									
2	Entry is equivalent to BAR at location 18h																									
3	Entry is equivalent to BAR at location 1Ch																									
4	Entry is equivalent to BAR at location 20h																									
5	Entry is equivalent to BAR at location 24h																									
6	Permitted to be used by a Function with a Type 1 Configuration Space header function only, optionally used to indicate a resource that is located behind the Function																									
7	Equivalent Not Indicated																									
8	Expansion ROM Base Address																									
9-14	Entry relates to VF BARs 0-5 respectively																									
15	Reserved – Software must treat values in this range as “Equivalent Not Indicated”																									
...																								
30	W	<p>Writable (W) – The value 1b indicates that the Base, and MaxOffset and Field Size fields for this entry are RW and that the Field Size bits for this entry are either RW or HwInit.; The value 0b indicates those fields are HwInit.</p> <p>See Table 6-x for additional requirements on the value of this field.</p>																								
31	E	<p>Enable (E) for this entry – 1b indicates this entry is enabled, 0b indicates this entry is disabled.</p> <p>If system software disables this entry, the resource indicated must still be associated with this function, and it is not permitted to reallocate this resource to any other entity.</p> <p>When system software writes or attempts to write to this bit, it must perform a read-modify-write such that all values read from all other bits of this DW are preserved.</p> <p>This field is permitted to be implemented as HwInit for functions that require the allocation of the associated resource, or as RW for functions that can allow system software to disable this resource, for example if BAR mechanisms are to be used instead of this resource.</p>																								

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

In the Enhanced Allocation ECN, page 12, make the following changes:

The Primary Properties[7:0] and Secondary Properties[7:0] fields are defined [in Table 7-34](#). ~~the following table:~~ [This table also defines whether or not the entry is permitted to be writeable. The Writeable bit in any entry must be 0b unless both the Primary and Secondary properties of that entry allow otherwise.](#)

Table 7-34: Enhanced Allocation Entry Field Value Definitions for both the Primary Properties and Secondary Properties Fields

Value (h)	Resource Definition	Writeable permitted
00	Memory Space, Non-Prefetchable. The corresponding Base and MaxOffset fields are HwInit.	No
01	Memory Space, Prefetchable. The corresponding Base and MaxOffset fields are HwInit.	No
02	I/O Space. The corresponding Base and MaxOffset fields are HwInit	No
03	For use only by Physical Functions to indicate resources for Virtual Function use, Memory Space, Prefetchable. The corresponding Base and MaxOffset fields are HwInit.	No
04	For use only by Physical Functions to indicate resources for Virtual Function use, Memory Space, Non-Prefetchable. The corresponding Base and MaxOffset fields are HwInit.	No
05	For use only by Type 1 Functions to indicate Memory, Non-Prefetchable, for Allocation Behind that Bridge. The corresponding Base and MaxOffset fields are HwInit.	No
06	For use only by Type 1 Functions to indicate Memory, Prefetchable, for Allocation Behind that Bridge. The corresponding Base and MaxOffset fields are HwInit.	No
07	For use only by Type 1 Functions to indicate I/O Space for Allocation Behind that Bridge. The corresponding Base and MaxOffset fields are HwInit.	No
08-FC	Reserved for future use; System firmware/software must not write to this entry, and must not attempt to interpret this entry or to use this resource. When software reads a Primary Properties value that is within this range, is it strongly recommended that software treat this resource according to the value in the Secondary Properties field, if that field contains a non-reserved value.	Yes

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

FD	Memory Space Resource Unavailable For Use - – System firmware/software must not write to this entry, and must not attempt to use the resource described by this entry for any purpose. The corresponding Base and MaxOffset fields are HwInit.	No
FE	I/O Space Resource Unavailable For Use - – System firmware/software must not write to this entry, and must not attempt to use the resource described by this entry for any purpose The corresponding Base and MaxOffset fields are HwInit.	No
FF	Entry Unavailable For Use – System firmware/software must not write to this entry, and must not attempt to interpret this entry as indicating any resource. It is strongly recommended that hardware use this value in the Secondary Properties field to indicate that for proper operation, the hardware requires the use of the resource definition indicated in the Primary Properties field .	No

In the Enhanced Allocation ECN, Figures 6-5, 6-6, 6-7, 6-8, and 6-9:

- Change “Reserved” for bits 39:24 of the first DW to “RsvdP”, and
- Change “R” for bit 3 of the first DW to “RsvdP”

B284 *VF Resizable Bar:* VF BAR Size

In the VF Resizable BARs ECN, page 7, Table 7-x3, make the following changes:

Bit Location	Description	Attributes
13:8	<p>VF BAR Size – This is an encoded value.</p> <p>0 = 1 MB (2^{20}) 1 = 2 MB (2^{21}) 2 = 4 MB (2^{22}) 3 = 8 MB (2^{23}) ... 43 = 8 EB (2^{63})</p> <p>The default value of this field is equal to the default size of the address space that the VF BAR resource is requesting via the VF BAR’s read-only bits.</p>	RW

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**


	<p>Behavior is undefined if a value is written in this field and the corresponding supported size bit is not Set in the VF Resizable BAR Capability or VF Resizable BAR Control registers.</p> <p><u>Software must only write values that correspond to those indicated as supported in the VF Resizable BAR Capability and Control registers. Writing an unsupported value will produce undefined results.</u></p> <p>When this register field is programmed, the value is immediately reflected in the size of the resource, as encoded in the number of read-only bits in the VF BAR.</p>	
--	---	--

B285 PASID TLP Prefix

In Section 6.20.2, page 629, Line 17, replace Figure 6-20 and update Table 6-13 as follows:

6.20.2. PASID TLP Layout

A TLP may contain at most one PASID TLP Prefix.



~~Figure 6-20. PASID TLP Prefix~~




Figure 6-20. PASID TLP Prefix

Table 6-13. PASID TLP Prefix

Bits	Description
31:29 <u>Byte 0</u> <u>Bits</u> <u>7:5</u>	100b – indicating TLP Prefix

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

28 Byte 0 bit 4	1b – indicating End-End TLP Prefix
27:24 Byte 0 Bits 3:0	0001b – indicating PASID TLP Prefix
23 Byte 1 Bit 7	Privileged Mode Requested – If Set indicates a Privileged Mode entity in the Endpoint is issuing the Request, If Clear, indicates a Non-Privileged Mode entity in the Endpoint is issuing the Request. Usage of this bit is specified in Section 6.20.2.3.
22 Byte 1 Bit 6	Execute Requested – If Set indicates the Endpoint is requesting Execute Permission. If Clear, indicates the Endpoint is not requesting Execute Permission. Usage of this bit is specified in Section 6.20.2.2.
21:20	Reserved
19:0 Byte 1 Bit 3: Byte 3 Bit 0	Process Address Space ID (PASID) – This field contains the PASID value associated with the TLP. Usage of this field is defined in Section 6.20.2.1.

B286 Multi-Function Arbitration Model Example

In Figure 6-10, Section 6.3.3.4, page 550, change Function 1 flow labeled “TC[2:4]” to “TC0, TC[2:4]”

and change “Egress Port” to “External Port”

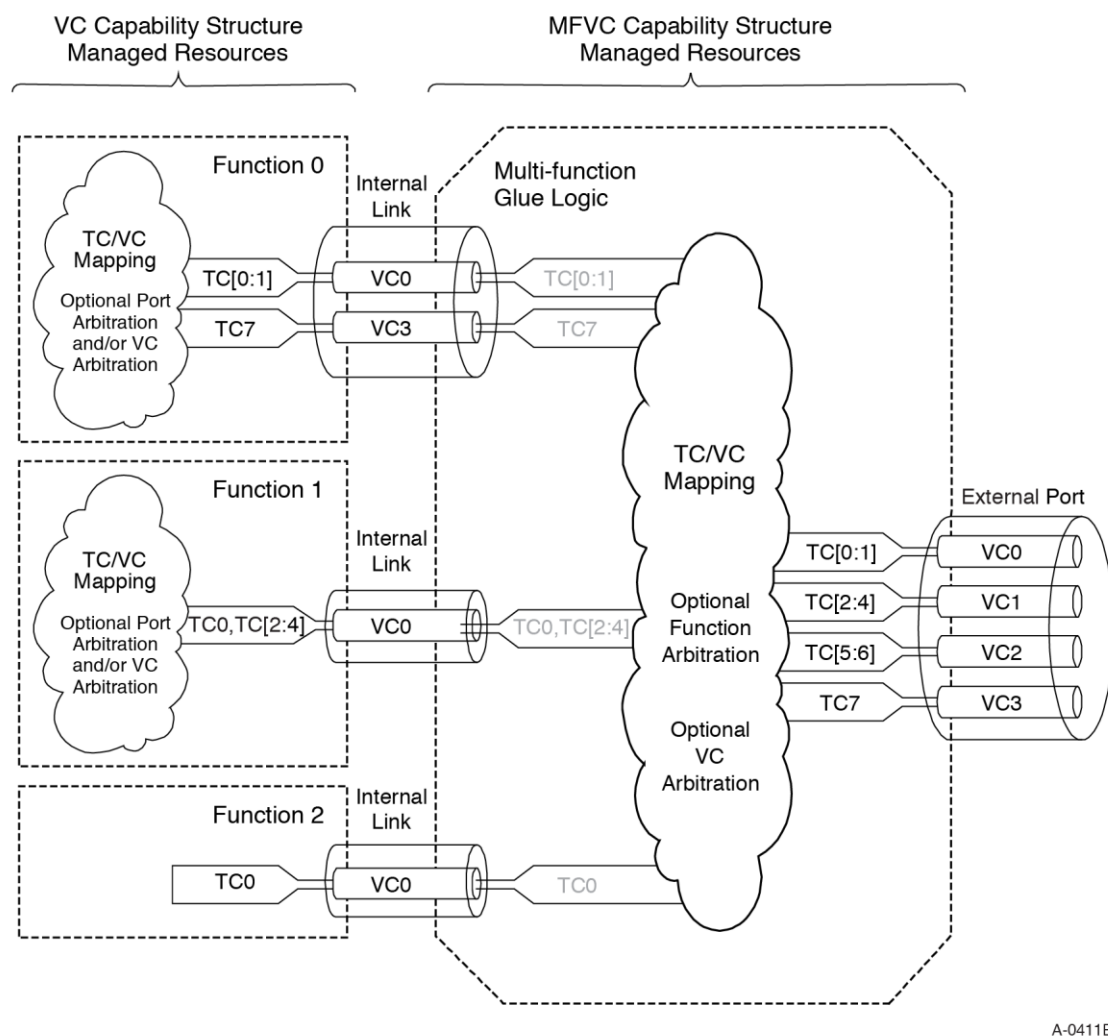


Figure 6-10: Multi-Function Arbitration Model

B288 Flow Control for Nullified TLPs

In Section 2.6, page 152, line 28 make the following changes:

2.6 Ordering and Receive Buffer Flow Control

...

Flow Control is handled by the Transaction Layer in cooperation with the Data Link Layer. The Transaction Layer performs Flow Control accounting functions for Received TLPs, and “gates”

TLP Transmissions based on available credits for transmission [even if those TLPs are eventually nullified](#).

B289 Clarification of HwInit

In Table 7-22, page 664 HwInit row, make the following changes:

Register Attribute	Description
HwInit	Hardware Initialized – Register bits are initialized by firmware or hardware mechanisms such as pin strapping or serial EEPROM. (System firmware hardware initialization is only allowed for system-integrated devices.) Bits are fixed in value and read-only after initialization and can only change in value or be reset (for write-once by firmware) with Fundamental Reset (see Section 6.6.1). HwInit register bits are not modified by an FLR.

B292 Capturing Device Number

In Section 2.2.6.2, page 83, line 2 make the following changes:

2.2.6.2 Transaction Descriptor – Transaction ID Field

...

- ❑ Functions must capture the Bus and Device Numbers⁸ supplied with all Type 0 Configuration Write Requests completed by the Function and supply these numbers in the Bus and Device Number fields of the Requester ID⁹ for all Requests initiated by the Device/Function. [It is recommended that Numbers are captured for successfully completed Requests only.](#)

Exception: The assignment of Bus and Device Numbers to the Devices within a Root Complex, and Device Numbers to the Downstream Ports within a Switch, may be done in an implementation specific way.

Note that the Bus Number and Device Number¹⁰ may be changed at run time, and so it is necessary to re-capture this information with each and every Configuration Write Request.

[It is recommended that Configuration Write Requests addressed to unimplemented Functions not affect captured Bus and Device Numbers.](#)

...

Note: Footnotes are not affected.

B293 DPC Header Log

In Section 7.31.11, page 893, line 2 make the following changes:

7.31.11. RP PIO Header Log Register (Offset 20h)

This register is implemented only in Root Ports that support RP Extensions for DPC. The RP PIO Header Log register contains the header from the [Request](#) TLP associated with a recorded RP PIO error. Refer to Section 6.2.10.3 for further details. This register is 16 bytes and is formatted identically to the Header Log register in AER. See Section 7.10.8.

B294 Completion Handling Rules for ID Based Ordering

In Section 2.3.2, page 138, line 38 make the following changes:

2.3.2 Completion Handling Rules

- ❑ When a device receives a Completion which does not match the Transaction ID for any of the outstanding Requests issued by that device, the Completion is called an “Unexpected Completion.”
- ❑ If a received Completion matches the Transaction ID of an outstanding Request, but in some other way does not match the corresponding Request (e.g., a problem with Attributes, Traffic Class, Byte Count, Lower Address, etc.), it is strongly recommended for the Receiver to handle the Completion as a Malformed TLP.
 - [The Completer must not check the IDO Attribute \(Attribute Bit 2\) in the Completion, since the Requester is not required to copy the value of IDO from the Request into the Completion for that request as stated in Section 2.2.6.4 and Section 2.2.9.](#)
 - However, if the Completion is otherwise properly formed, it is permitted²⁴ for the Receiver to handle the Completion as an Unexpected Completion.

B295 Downstream Port Containment Clarifications

In Section 6.2.10, page 531, line 16 and line 22, make the following changes:

6.2.10. Downstream Port Containment (DPC)

...

When DPC is triggered, the Downstream Port immediately Sets the DPC Trigger Status bit and DPC Trigger Reason field to indicate the triggering condition (unmasked uncorrectable error, ERR_NONFATAL, ERR_FATAL, RP_PIO error, or software triggered), and disables its Link by directing the LTSSM to the Disabled state. Once the LTSSM reaches the Disabled state, it remains in that state until the DPC Trigger Status bit is Cleared. To ensure that the LTSSM has time to reach

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

the Disabled state or at least to bring the Link down under a variety of error conditions, software must leave the Downstream Port in DPC until the Data Link Layer Link Active bit in the Link Status register reads 0b; otherwise, the result is undefined. See Section 7.8.8. See Section 2.9.3 for other important details on Transaction Layer behavior during DPC.

...

After software releases the Downstream Port from DPC, the ~~associated Link will normally~~ Port's LTSSM must transition to the Detect state, where the Link will attempt to retrain. Software can use Data Link Layer State Changed interrupts, DL_Active ERR_COR signaling, or both, to signal when the Link reaches the DL_Active state again. See Sections 6.7.3.3 and 6.2.10.5.

In Table 7-16, page 708, make the following changes:

Table 7-16: Link Control Register

Bit Location	Register Description	Attributes
...
5	<p>Retrain Link – A write of 1b to this bit initiates Link retraining by directing the Physical Layer LTSSM to the Recovery state. If the LTSSM is already in Recovery or Configuration, re-entering Recovery is permitted but not required. <u>If the Port is in DPC when a write of 1b to this bit occurs, the result is undefined</u>. Reads of this bit always return 0b.</p> <p>It is permitted to write 1b to this bit while simultaneously writing modified values to other fields in this register. If the LTSSM is not already in Recovery or Configuration, the resulting Link training must use the modified values. If the LTSSM is already in Recovery or Configuration, the modified values are not required to affect the Link training that's already in progress.</p> <p>This bit is not applicable and is Reserved for Endpoints, PCI Express to PCI/PCI-X bridges, and Upstream Ports of Switches.</p> <p>This bit always returns 0b when read.</p>	RW

In Table 7-134, page 885, make the following changes:

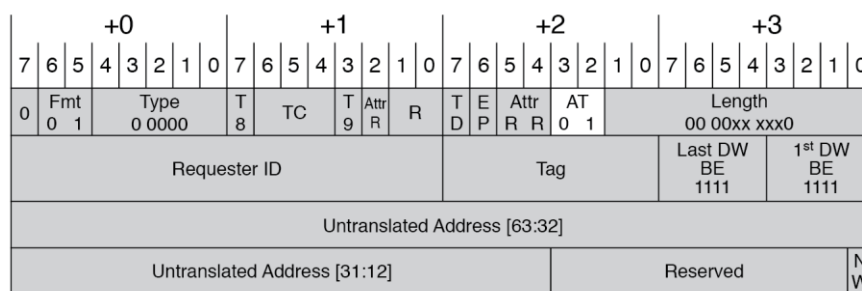
Table 7-134. DPC Status Register

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

Bit Location	Register Description	Attributes
0	<p>DPC Trigger Status – When Set, this bit indicates that DPC has been triggered, <u>and by definition the Port is “in DPC”</u>. DPC is event triggered.</p> <p>While this bit is Set, hardware must direct the LTSSM to the Disabled State. This bit must be cleared before the LTSSM can be released from the Disabled State, <u>after which the Port is no longer in DPC, and the LTSSM must transition to the Detect State</u>. See Section 6.2.10 for requirements on how long software must leave the Downstream Port in DPC. Once these requirements are met, software is permitted to clear this bit regardless of the state of other status bits associated with the triggering event.</p> <p>After clearing this bit, software must honor timing requirements defined in Section 6.6.1 with respect to the first Configuration Read following a Conventional Reset.</p> <p>Default value of this bit is 0b.</p>	RW1CS

B297 *ATS: Clarifications*

In ATS 1.1, page 22, Figure 2-4 change the length field 00 000xxx xxx0




A-0578C

Figure 10-8 64-bit Translation Request Header

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

In ATS 1.1, page 23, Figure 2-3 change the length field 00 00~~0xx~~ ~~xxx~~0



A-0621C

Figure 10-9: 32-bit Translation Request Header

In ATS 1.1, Section 5.1, Page 45 make the following changes:

5.1 ATS Extended Capability Structure

Each Function that supports ATS ([capable of generating Translation Requests](#)) must have the ATS Extended Capability structure in its extended configuration space. [It is permitted to be implemented by Endpoints or Root Complex Integrated Endpoints.](#)

...

In ATS 1.1, Section 5.2, page 47, make the following changes”

5.2 Page Request Extended Capability Structure

A Page Request Extended Capability Structure is used to configure the Page Request Interface 5 mechanism. A Multi-Function [Endpoint or Root Complex Integrated Endpoint](#) Device may implement a Page Request Interface and the associated capability on any Function [within the Device](#). For SR-IOV, a single Page Request Interface is permitted for the PF and is shared between the PF and the associated VFs. The PF implements this capability and the VFs do not. Every Page Request Interface mechanism operates independently.

B299 **PCI Local Bus and SR-IOV: Vendor, Device, and Revision IDs**

In Section 6.2.1, page 216 of the PCI Local Bus Specification, make the following changes:

6.2.1 Device Identification

...

Vendor ID This field identifies the manufacturer of the device. [In keeping with PCI-SIG procedures](#), ~~V~~valid vendor identifiers ~~are~~ must be obtained from the PCI-SIG to ensure uniqueness. [Each vendor must have at least one Vendor ID](#). FFFFh is an invalid value for *Vendor ID*. [It is recommended that software read the Vendor ID register to determine](#)

if a Function is present, where a value of FFFFh indicates that no Function is present.

Device ID This field identifies the particular device. ~~This identifier is~~ The Device ID must be allocated and tracked by the vendor. The Device ID, in conjunction with the Vendor ID and Revision ID, are used as one mechanism for software to determine which driver should be loaded. The vendor must ensure that the chosen values do not result in the use of an incompatible device driver. ~~This field should be viewed as a vendor defined extension to the Device ID.~~

Revision ID This register specifies a device specific revision identifier. The value is chosen by the vendor. Zero is an acceptable value. The Device ID, in conjunction with the Vendor ID and Revision ID, are used as one mechanism for software to determine which driver should be loaded. The vendor must ensure that the chosen values do not result in the use of an incompatible device driver.

...

In Section 6.2.4, page 224 of the PCI Local Bus Specification, make the following changes:

6.2.4 Miscellaneous Registers

...

Subsystem Vendor ID and Subsystem ID

~~These~~ The Subsystem Vendor ID and Subsystem ID registers are used to uniquely identify the add-in card, motherboard, or subsystem where the PCI device resides. They provide a mechanism for add-in card or subsystem vendors to distinguish their add-in cards or subsystem from one another even though the add-in cards may have the same PCI controller on them (and, therefore, the same Vendor ID and Device ID).

Implementation of these registers is required for all PCI devices except those that have a ~~b~~Base ~~e~~EClass 06h with ~~s~~Sub-eClass 00h to 04h (00h, 01h, 02h, 03h, 04h), or a ~~b~~Base ~~e~~EClass 08h with ~~s~~Sub-eClass 00h to 03h (00h, 01h, 02h, 03h). Subsystem Vendor IDs can be obtained from the PCI SIG and are used to identify the vendor of the add-in card, motherboard, or subsystem.⁴⁴ A Subsystem Vendor ID (SVID) must be a Vendor ID assigned by the PCI-SIG to the vendor of the subsystem. Subsystem Vendors must obtain a Vendor ID from the PCI-SIG. In keeping with PCI-SIG procedures, valid vendor identifiers must be obtained from the PCI-SIG to ensure uniqueness. ~~The SVID must match the Vendor ID of the vendor of the subsystem.~~

Values for the Subsystem ID are vendor assigned. Subsystem ID values, in conjunction with the Subsystem Vendor ID, form a unique identifier for the PCI product. Subsystem ID and Device ID values are distinct and unrelated to each other, and software should not assume any relationship between them. Values in these registers must be loaded and valid prior to the system firmware or any system software accessing the PCI Configuration Space. How these values are loaded is not specified but could be done during the manufacturing process or loaded from external logic (e.g., strapping options, serial ROMs, etc.). These values must not be loaded using expansion ROM software because expansion ROM software is not guaranteed to be run during POST in all systems.

Errata for the PCI Express Base Specification Rev. 3.1a, SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0

Devices are responsible for guaranteeing the data is valid before allowing reads to these registers to complete. ~~This can be done by responding to any accesses with Retry until the data is valid~~

If a device is designed to be used exclusively on the system board, the system vendor may use system specific software to initialize these registers after each power-on.



IMPLEMENTATION NOTE

Subsystem Vendor ID and Subsystem ID

The Subsystem Vendor ID and Subsystem ID fields, taken together, allow software to uniquely identify a PCI circuit board product. Vendors should therefore not reuse Subsystem ID values across multiple product types that share a common Subsystem Vendor ID. It is acceptable, although not preferred, to reuse the Subsystem ID value on products with the same Subsystem Vendor ID in cases where the products are in the same family and generation, and differ only in capacity or performance. Note also that it is acceptable for vendors to use multiple unique Subsystem ID values over time for a single product type, such as to indicate some internal difference such as component selection.

In the SR-IOV Specification, page 41, line 1 make the following changes:

3.3.11 VF Device ID (1Ah)

This field contains the Device ID that should be presented for every VF to the SI.

VF Device ID may be different from the PF Device ID. A VF Device ID must be tracked by the vendor. It is recommended that the vendor must ensure that the chosen VF Device ID does not result in the use of an incompatible device driver.

B300 PME Generation

In Section 5.3.1.4, Page 449, line 28, make the following changes:

5.3.1.4. D3 State

D3 support is required, (both the D3cold and the D3hot states). ~~Functions supporting PME-generation from D3 must support it for both D3cold and the D3hot state.~~

**Errata for the PCI Express Base Specification Rev. 3.1a,
SR-IOV Rev. 1.1, ATS Rev. 1.1, and M.2 Rev. 1.0**

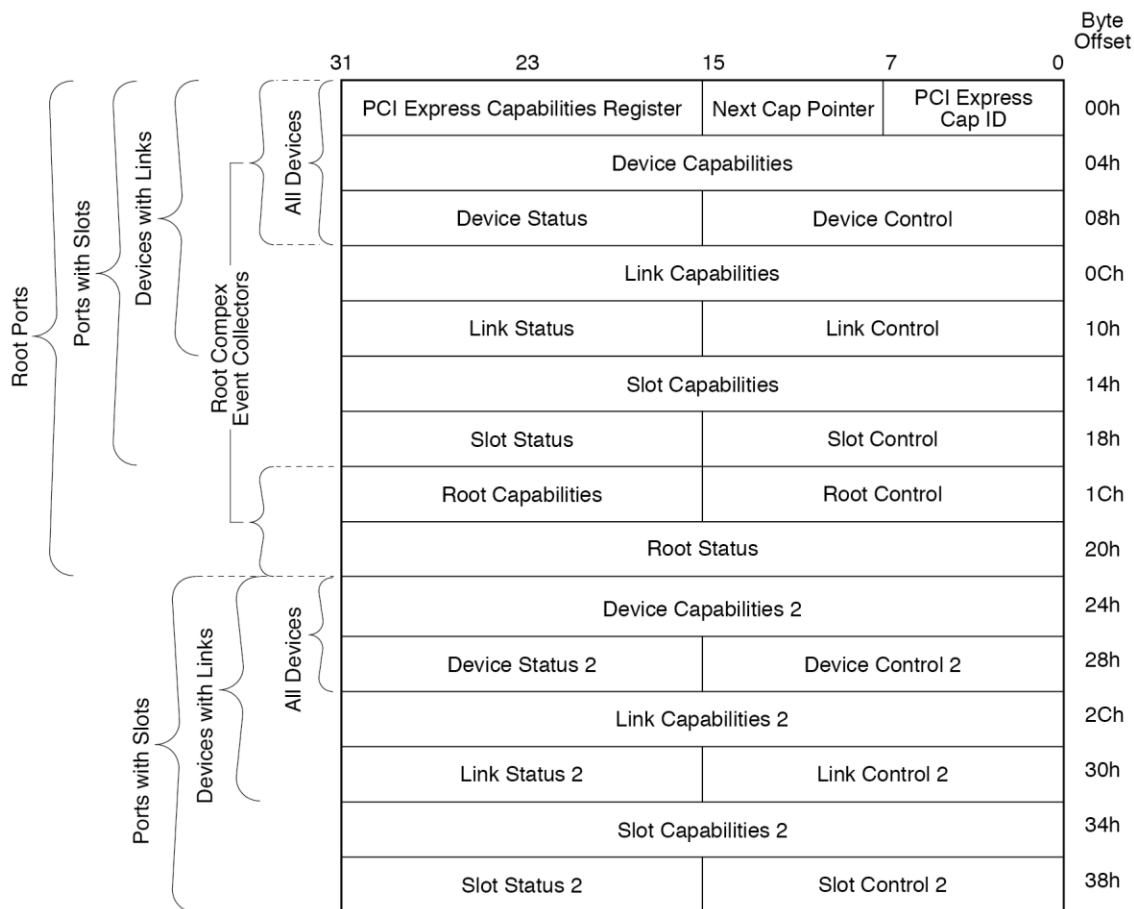
In Section 7.6, Table 7-8, Page 679 make the following changes:

Table 7-8: Power Management Capabilities Register Added Requirements

Bit Location	Register Description	Attributes
31:27	<p>PME Support – For a device Function, this 5-bit field indicates the power states in which the Function may generate a PME and/or forward PME Messages.</p> <p>Each bit that corresponds to a supported D-state Bits 31, 30, and 27 must be Set for PCI-PCI Bridge structures representing Ports on Root Complexes/Switches to indicate that the Bridge will forward PME Messages. Bit 31 must only be Set if the Port is still able to forward PME Messages when main power is not available.</p>	Unchanged

B301 PCI Express Capability Structure

In Figure 7-10, Page 683, add the 3 braces shown at the left (the text was present in Base 3.1a but the braces were missing):



OM14318B

Figure 7-10: PCI Express Capability Structure

B303 **Hierarchy ID ECN: Hierarchy ID Message**

In the Hierarchy ID ECN, page 4, Figure 2-xx:

- Change the “Type” field to ~~01~~10 011
- Change the “Tag” field at byte 5 to “Reserved”

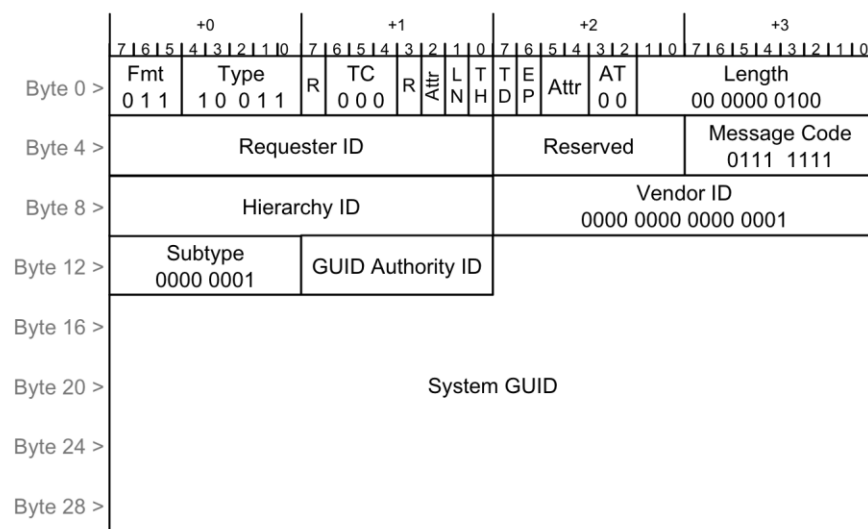


Figure 2-xx Hierarchy ID Message